FlexiVol: a Volumetric Display with an Elastic Diffuser to Enable Reach-Through Interaction

Elodie Bouzbib
UpnaLab
Universidad Publica de Navarra
Pamplona, Spain
elodie.bouzbib@unavarra.es

Iosune Sarasate
Universidad Pública de Navarra
Pamplona, Spain
iosune.sarasate@unavarra.es

Unai Javier Fernández
Universidad Pública de Navarra
Pamplona, Spain
unai.fernandez@unavarra.es

Ivan Fernández
UpnaLab
Universidad Pública de Navarra
Pamplona, Navarra, Spain
ivan.fernandezf@unavarra.es

Manuel Lopez-Amo
Universidad Pública de Navarra
Pamplona, Spain
manuel.lopezamo@unavarra.es

Iñigo Ezcurdia
Public University of Navarra
Pamplona, Spain
inigofermin.ezcurdia@unavarra.es

Asier Marzo
UpnaLab
Universidad Publica de Navarra
Pamplona, Navarre, Spain
asier.marzo@unavarra.es

Figure 1: FlexiVol is a volumetric display with an elastic diffuser that allows users to reach inside and interact directly with true 3D graphics. A) The user's hand and the virtual objects have coherent focal accommodation: 1. the view is focused at the hand and red cube, or 2. at the sphere behind them. B) The user input (hand) is aligned with the output (rendered objects), allowing to directly 1. grab virtual objects, 2. to translate and 3. rotate them.

Abstract

Volumetric displays render true 3D graphics without forcing users to wear headsets or glasses. However, the optical diffusers that volumetric displays employ are rigid and thus do not allow for direct interaction. FlexiVol employs elastic diffusers to allow users to reach inside the display volume to have direct interaction with true 3D content. We explored various diffuser materials in terms of visual and mechanical properties. We correct the distortions of the volumetric graphics projected on elastic oscillating diffusers and propose a design space for FlexiVol, enabling various gestures and actions through direct interaction techniques. A user study suggests that selection, docking and tracing tasks can be performed faster and more precisely using direct interaction when compared to indirect interaction with a 3D mouse. Finally, applications such
as a virtual pet or landscape edition highlight the advantages of a volumetric display that supports direct interaction.

CCS Concepts

• **Hardware** → *Emerging interfaces*; • **Human-centered computing** → **Visualization systems and tools**; • **Computing methodologies** → **Graphics systems and interfaces**.

Keywords

Volumetric Displays, Direct Interaction, True 3D graphics, Flexible Diffuser, Projection

1 Introduction

We are used to view and interact in a 3D world. Our vision uses cues to interpret the space, and with our hands we directly grab and manipulate the objects within.

Graphics rendered on a 2D screen can provide monocular visual cues such as occlusion, distance-size relationship, shadows or texture gradients [16] but binocular disparity, accommodation of the focal point or convergence cannot be displayed. Differently, a true 3D display [30] renders graphics that can be viewed by multiple people from different angles and provide those visual cues; also, they do not force the users to wear any device.

Direct interaction is a "natural" way of interacting with virtual entities in which the output space (rendered graphics) is aligned with the input space (interaction area) [21] – similar to our interactions with the real world objects [6]. This type of interaction is widely employed in multi-touch flat screens: we directly press on the buttons with the finger, drag icons to move them, or rotate an object with two fingers as we would do to rotate a sheet of paper in real life.

Volumetric displays render true 3D and can provide most of the visual cues that we perceive from the real world. However, with current technologies, users cannot reach inside the display volume to directly interact with virtual objects as they would do in real life. Their diffusers are most often rigid, which causes safety hazards when reaching through; and interaction techniques are thus performed indirectly using a 3D mouse or a keyboard.

FlexiVol's concept is based on modifying volumetric displays with an elastic optical diffuser to enable users to insert their hand within the rendering volume and directly interact with spatially overlapped true 3D graphics. The true 3D graphics and user hand provide coherent focus accommodation (Figure 1.A), which allows for enhanced depth perception. The interaction with the rendered graphics is direct, the user can grab an object to move and rotate it (Figure 1.B).

We demonstrate FlexiVol's concept in a modified commercially available volumetric display (Voxon VX1) and in custom-made volumetric displays. We studied different candidates for the diffuser material, and proposed parameters to guide their selection as elastic diffusers: elastic properties, optical diffusion, frequency response and amplitude. We also proposed a simple method to correct the visual distortion introduced by the oscillating elastic diffuser.

Inspired by interactions with 2D elastic diffusers [57], we propose a design space of interaction techniques for elastic diffusers in 3D volumetric displays, through users gestures and actions *within* FlexiVol. We then conducted a user study (n = 18), consisting of three tasks (Selection, Tracing, Docking) to validate FlexiVol and compare the reach-through input modality with a 3D mouse, a common input modality for current volumetric displays.

We show that (1) reach-through interaction techniques significantly improve completion time, (2) self-assessed performance and (3) cognitive workload; (4) as a trade-off with accuracy. Qualitative feedback also (5) indicates how reaching-through FlexiVol to select and manipulate 3D objects may initially seem intimidating or bothersome, but users quickly discover the interaction to be safe and surprisingly gentle to the touch.

The contributions of this paper are:

- FlexiVol, a novel concept of using elastic diffusers for volumetric displays,
- Guidelines for characterizing and selecting elastic diffusers for volumetric displays,
- A compatible design space for reach-through interaction techniques,
- Empirical results of completion time, accuracy and precision and mental workload comparing reach-through interaction with a 3D mouse,
- Proposed applications using FlexiVol.

2 Related Work

2.1 True 3D and Volumetric Displays

In a true 3D display [30], graphics can be observed by multiple people from different points of view, without forcing them to wear any device, and with the depth visual cues that we get from the real world.

Head-Mounted-Displays (HMDs) provide depth cues by showing different images to each eye thus creating binocular disparity but they usually do not provide convergence and focus accommodation [52], thus users cannot focus correctly at their hands and nearby objects. State of the art HMDs are exploring eye-tracking [2] or holographic near-field displays [36, 41] to support convergence or focal accommodation. However, even if those technologies get fully developed, each user still has to wear a head-mounted display, hindering the come and interact paradigm where one or various users just approach a system and start using it.

From the different 3D technologies, only volumetric displays and holograms provide all the depth cues [47]. A volumetric display [19] emits points of light from each position within a volume, this technology is superior to holograms which present problems such as clipping and forbidden geometries [53].

Volumetric displays have been classified as: swept, solid and free-space [54]. Swept volumetric displays [51] exploit the persistence of vision to render different fractions of a 3D scene at different positions, so that they are perceived as a single three-dimensional scene [3, 28]. They use high-speed moving rigid parts which can be dangerous to touch. Solid displays use non-linear optical media that emits light when illuminated by two different wavelengths coming from perpendicular directions [42], or two-photon absorption [35]. Free-space display examples include plasma induced in mid-air using a powerful focused laser [50]; they involve severe dangers. Other free-space volumetric displays use levitated tracer particles trapped with optical [53] or acoustic tweezers [15, 22] .

In those volumetric displays, the user cannot insert the hand: it can damage the user or the display, and it is not physically possible to touch a solid cube or the levitation stops working. In other words, these existing true 3D displays do not allow the users to reach inside the display volume to directly interact with the rendered objects, removing the possibility to combine true 3D graphics with direct interaction. Plasma induced from lasers can be made touchable by femtosecond pulsing [40] or optically redirecting the graphics using parabolic mirrors [32]; however, these alternatives are dangerous to the eye and only produce rendering volumes below 2 cm.

2.2 Interaction with Volumetric displays

Several interaction techniques have been developed for 3D environments [10, 26], but mostly indirect approaches have been adapted for volumetric displays [17] such as rays [13] or cursors. Other alternatives use hand gestures [4] or the index finger as a pointing device [18]. Methods based on touching the display boundary (i.e., protective case) [29] try to get closer to the display volume, yet the interaction remains indirect.

Commercial volumetric displays such as the Voxon VX1 [44] use traditional mouse and keyboard techniques, or 3D mice [1] as an input modality. Voxon displays also support a hand-tracking device (Leap Motion, Ultraleap Ltd.) enabling the use of the hands to manipulate virtual objects, but the real hands have to remain outside of the display volume.

In this paper, we investigate how to enable *reach-through interactions* to facilitate the selection and manipulation of true 3D content inspired from the definition of *direct interaction* from [21]: *allowing the user to literally get their hands into the virtual display and to directly interact with a spatially aligned 3D virtual world.* We propose the use of non-rigid diffusers on swept volumetric displays to enable the user to reach inside the display volume.

2.3 Non-rigid Diffusers

Fog screens enable users to pass their hands through 2D projection screens [45]. Shape-changing fog-screens [56] can provide a 2.5D surface for projection but the control points of the surface are constrained (5 control points along the vertical direction; i.e., 1D) and volumetric graphics are not possible. A matrix of 6×6 fog emitters allows flat projections to appear at discrete positions [33]. Water droplets can also form a projection screen, allowing to stack up to 3 droplet screens at different depths [5]. A 2D array of falling droplets can act as a diffuser for volumetric graphics [11]; however, the rendered graphics were spheres or cubes occupying the whole display due to limited resolution along the falling direction. None of these approaches provide continuous volumetric graphics or with enough resolution to study direct interaction techniques. Furthermore, no user study or exploration of the direct interaction capabilities was conducted. HoloDust [43] is a theoretical concept that projects in falling dust taking into account its particle distribution to render true 3d graphics, but was never implemented.

An alternative to obtain continuous graphics with direct interaction, is to replace a solid diffuser (e.g., tablet screen) by an elastic one (e.g., a fabric) to enable new interaction techniques such as pushing or pinching it [8, 31, 39, 49, 58, 59]. However, these displays only have 2D or limited 2.5D [9] rendering capabilities. In this paper, we draw inspiration from the use of elastic diffusers and their taxonomy of interaction techniques [57] to enable novel interaction techniques *within* a volumetric display with elastic diffusers.

3 Swept Volumetric Display System Design

We relied on three different mechanisms to realize and explore FlexiVol's concept. A modified commercially available swept volumetric display (Voxon VX1) (Figure 2 - C) and two custom-made devices for our technical explorations (Figure 2 - A,B) that enabled us to control parameters such as oscillation frequency and amplitude as well as projection patterns. These mechanisms rely on the same

principle, described in this section. A screenshot of a custom-made swept volumetric display is available in Appendix A.1; and files for replication are available here.

Swept volumetric displays render true 3D content by sweeping a 2D surface, its "diffuser", at high speed. The graphics are projected upon this surface, as 2D slices, for each position of the diffuser a different slice is projected, persistence of vision makes the users perceive it as a 3D volume (see Figure 3). The diffuser usually oscillates following a sinusoid, the oscillation is generated by one or more actuators (such as a voicecoils) and its displacement can be amplified by a mechanical amplifier (e.g., planar spring).

Figure 2: (A, B, C) Schematics of three swept volumetric displays with elastic diffusers. D) Actuators input signal, frame oscillation, and diffuser movement.

3.1 Hardware

Projector: A high frame rate is needed to render into the display volume: if 100 slices are projected at a frame rate of 30 Hz, the projector has to operate at 3000 Hz. Usually, DLP projectors with time multiplexation for the color are employed. The projected images are reduced to binary patterns and thus the original frame rate of the projector can be multiplied by the number of color bitplanes. For example, the Voxon VX1 projects 192 slices at 15 Hz (2880 Hz) with a resolution of 980×720px. In our custom-made volumetric display with *one actuator* (see Figure 2 - A or Appendix A.1), we used a DLP LightCrafter 4500. When operating at 120 Hz with 24-bit color depth, it projects binary patterns at 2880 Hz with a resolution of of up to 1280×800px. Color volumetric rendering is possible by losing depth resolution and multiplexing in space the color (e.g, odd slices are projected blue whereas even slices are red).

Actuators: the typical actuators for a planar oscillating SVD are voicecoils (e.g., Voxon VX1 [44] and Dan Maloney's [37]) either custom designed or from a modified speaker. The VVD [24] uses stepper motors that rotate back and forth 25° as an alternative. In our custom-made volumetric displays (see Figure 2 - A,B), we employed a modified VDSSP6.5/8 HQ Power speaker.

Mechanical amplifier: the actuators may not oscillate with sufficient amplitude for the diffuser to provide a large enough display volume. That is why some designs (e.g., Voxon VX1 and Dan Maloney's) use mechanical amplifiers. A mechanical amplifier can be just a spring or a lever mechanism. In our custom-made *one actuator* volumetric display (see Figure 2 - A), we connected the speaker to an audio amplifier (Fosi Audio TP-02); and used a mechanical PLA 3D-printed spring (3 mm thickness, 70 mm diameter, 3 spirals). The design with *two-actuators* does not use mechanical amplifiers to only study the mechanical behaviour of the diffusers.

Diffuser: it is usually a sheet of an optically diffusing material, which receives the projected light and diffuses it towards the viewers. The display volume is usually protected with an acrylic case to avoid users touching the solid diffuser. For FlexiVol, we replaced this traditional solid diffuser with an array of elastic strips (see Section 4).

3.2 Software

Rendering Rendering for a volumetric display is different from typical scanline or ray tracing pipelines. The scene needs to be sliced along planes, those slices are projected at high-speed on the oscillating diffuser (see Figure 3). Therefore, several slices are combined into a color image to be sent to the projector. For example, 24 slices with binary color are packed into a regular 24-bit RGB image that the projector is expecting from the HDMI or Display Port input.

If the 3D scene is defined as voxels, either from a 3D-array or a function, the rendering is simplified since each pixel from the slice can be directly obtained from the corresponding voxel.

If the application uses typical meshes with vertices and triangles, the rendering of the slices is not trivial. Slicing algorithms for hollow rendering use a fragment shader that renders pixels as pure white if its 3D position is between the sliced depth or otherwise get discarded. For solid rendering of the meshes, the z-clipping minimum value is set to the slice depth and only the back-facing triangles are rendered in pure white so that they fill in the cutout of that object.

Synchronization In general, even if the actuators are excited with the same frequency as the projector, there will be drift and synchronization is required (see Figure 3 - the slices of the 3D object are overlaid over the sinusoid of the actuators input).

If the actuator signal is used to synchronize the display, the rendering system has to sample the actuation signal before the start of each frame. That is, the rendering system checks the phase of the actuation signal and extrapolates the diffuser position for the next 24 slices, renders the slices and sends them packed as a color image to the projector.

If the projector is the synchronization source, the vertical sync signal can be used to keep in a phased-locked-loop the actuation signal. The actuation signal is required to be a divider of the vsync signal. For example, in the Voxon VX1 the actuation signal is a 15 Hz sinusoidal and the vsync of the projector is a 120 Hz signal.

4 Elastic Diffusers

For FlexiVol, we replace the rigid diffuser of a planar swept volumetric display by an elastic material that allows for deformations
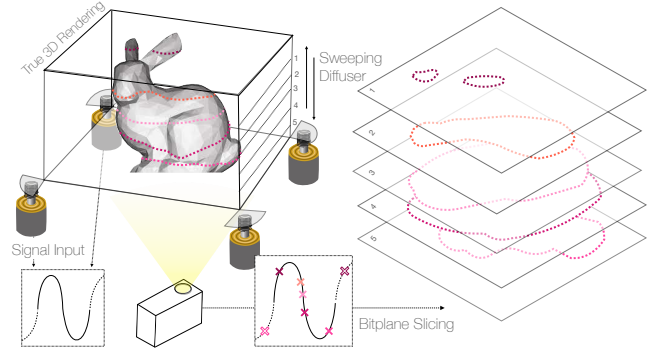
Figure 3: Rendering in a swept volumetric display. The actuators are excited with a sinusoidal signal; the projector is synchronized with this input to project different slices at specific positions of the sweeping diffuser. The slices are a cross-section of the 3D object; that, at high speed oscillation, are perceived as a whole volume.

without harming the user or damaging the display. The most common used materials for elastic projection screens are Lycra and latex [57], we thus used these materials as a starting point but also characterized others.

For the diffuser, we decided to use an array of strips rather than a continuous membrane because when the user pushes the membrane, it creates deformations along all the membrane and distorts the entire display volume; with the strips, only the pushed ones get distorted. Furthermore, the forces that a membrane exerts on the user's finger or the mechanical amplifiers can damage them, especially for anisotropic materials that have a large stiffness along their non-elastic dimension. Finally, the deformation of an oscillating 2D membrane is not as predictable as a single strip, the latter can just be approximated with a sinusoid or parabole (see Section 5). We employed strips of 20 mm width to approximately match the finger width; smaller strips had the same distortion upon user contact, produced more optical discontinuities and the arrays were more cumbersome to adjust. Larger strips tended to have the same issues as a full membrane: too much force and distorted area upon contact with the hand.

All the fabrics were laser-cut into $200 \times 20 \; mm^2$ strips along their elastic dimension. The curable materials (e.g., silicone) were cured on top of an acrylic sheet with a spacer of the desired thickness.

The selected materials samples can be seen in Figure 4 and are:

Elastane also called Lycra, Spandex or Dorlastan is a synthetic fiber (polyether-polyurea copolymer). We used fabrics of thicknesses 0.8 mm (elastane A), 0.5 mm (elastane B) and 0.8 mm with a 20% polyester (elastane C) - Figure 4 - 1 to 3.

Silicone we used (Silicone RPRO Glass, Reschimica) and cured them in sheets of thickness 0.8 mm - Figure 4 - 4.

Elastic band we used an elastic band typically used for clothes, for example along the waist of stretchable trousers. We picked a composition of 60% polyester and 40% elastane with a thickness of 1.0 mm - Figure 4 - 5.

Projection screen commercial fabric used for projection screens (120" double side, Osoeri) made of polyester and a thickness of 0.5 mm - Figure 4 - 6.
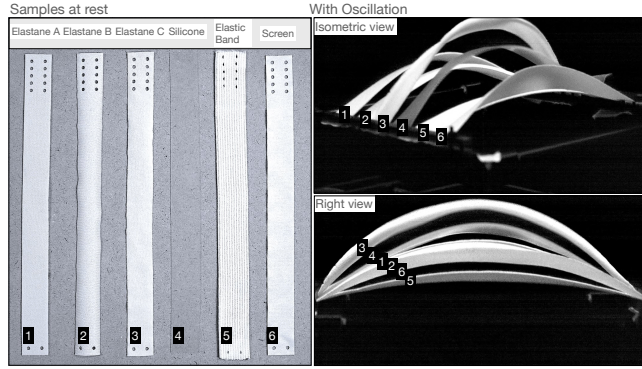

Figure 4: Left) Strip samples. Right) Strips oscillating (15 Hz).

Discarded materials: Organza, silk and other fabrics were too stiff, they applied too much force on the user finger and display. Latex (latex pro, prevulcanized, Reschimica) was also tested in 0.7 mm sheets but it offered no advantages compared to silicone [34].

We tried a composite material by infusing the projection screen with silicone, trying to obtain the visual quality of the former and elastic properties of the latter. However, the resulting strip was too stiff and prone to detachment between its components. There are fabrics specifically designed for back-projection but we found them to have too much stiffness and density.

4.1 Optical diffusion

To assess the optical diffusing properties of the strips, we employed a back-projection setup to project test images on the selected samples. These projections were performed with the samples in both unstretched and stretched states.

Images were captured using a Sony Alpha 6300 camera equipped with a Sony E PZ 16-50 mm F3.5-5.6 OSS lens, positioned 30 cm above the samples in a dark environment. The camera settings were: zoom at 21 mm, shutter speed 1/30 s, aperture F8.0, metering mode Multi-segment (M-M) with a value of 1.3, and ISO 100. An EZCast Beam J2 projector (native 854x480 resolution, 300 lumens, contrast 5000:1) was positioned 30 cm below the samples.

Test patterns, including blank white light, a grid on a black background, and RGB stripes, were projected to assess brightness, sharpness, color accuracy, or distortion. All patterns were projected at three elongations: the original length of 185 mm, a stretched length of 210 mm, and a maximum stretched length of 235 mm. Figure 5 presents results only for the original and maximum stretched lengths (185 mm and 235 mm, respectively) for clarity.

We chose a maximum elongation of 50 mm (from 185 to 235 mm) as it approximates the maximum elongation caused by the finger pushing down the strip center on a display volume of 190×190×80 mm^3. This also avoids irreversible plastic deformations on the samples (see Section 4.4).

A qualitative analysis of Figure 5.A (and Appendix A.2) reveals that, as expected, the *Projection Screen* provides even light diffusion,


Figure 5: Optical projection on the selected materials. A.Left) Back-projected patterns (RGB stripes, a black and white grid, and white light). A. Zenithal photographs with the strips at their original length of 185 mm (Middle) and stretched to 235 mm (Right). B) Red, green and blue average Savitzky-Golay filtered intensities at different offset angles from the central projection axis.

while *Silicone* exhibits a prominent central flare. *Elastane B* exhibits diffusion comparable to *Projection Screen*. Dark artifacts due to their weaving pattern are observed in *Elastane A*, *Elastane C*, and particularly in *Elastic Band* samples.

Inspection of the B&W grid pattern and the RGB stripes pattern reveals that *Elastane B* and *Elastane C* distort the projected graphics when they are stretched. In contrast, *Silicone* displays sharp and well-defined lines; however, the aforementioned glare significantly impacts certain areas, causing overexposure. The *Elastic Band* does not distort the graphics upon stretching, yet straight lines aligned with the thread direction may be prone to artifacts.

Figure 5.B presents red, green and blue intensities at different offset angles from the central projection axis for *Silicone*, *Elastic Band* and *Projection Screen materials* (comparative graphs for all materials at both original and stretched lengths are available in Appendix A.2).

We note that *Silicone* produces a bleached final image. In contrast, *Elastic Band* shows a more consistent behaviour across its surface; similarly to the traditional projection *Screen*. However, the *Elastic Band* affects the projection with dark stripes along its thread orientation while other materials, such as the *Projection Screen* or *Elastane A*, maintain uniform color.

4.2 Stiffness

We measured the transversal stiffness of the samples. That is, the force that the material exerts upwards as it is pushed down along its center. A 3D printer (Ender 3 Pro) was modified by substituting the hot-end with a load cell (1 kG HALJIA with an HX711 ADC). The samples were mounted on a frame on top of the printer bed and the load cell pressed down the strip. A computer controlled the printer with g-code and read the force-weight at each displacement position. The setup and results are shown in Figure 6.
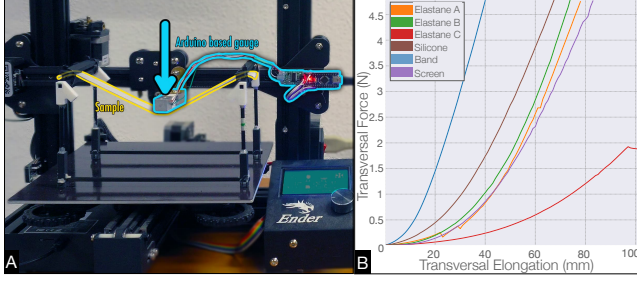
Figure 6: A. Setup for measuring transverse stiffness: A 3D printer positions a load cell to deform the sample. B. Transversal stiffness of the samples.
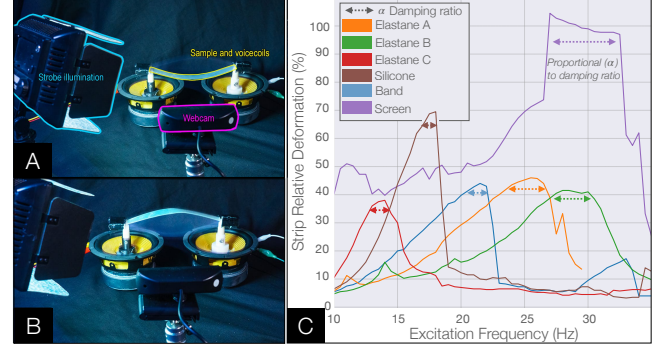
4.3 Frequency Response

We experimentally captured the frequency response of the strips. The oscillation amplitude at the center of the strip was measured for different frequencies of the excitation signal, ranging from 10 Hz to 30 Hz. We used a custom-made mechanism (shown in Figure 2 - B). We set the settling time between frequencies at 200 ms, and captured 200 frames at 25 FPS with a LOETAD 1080P HD Webcam. Strobe illumination was used at +0.5Hz from the excitation signal to allow the capture of the oscillation at a regular frame rate, since our available high-speed cameras were not able to store the whole frequency sweeping process, this setup allowed to capture phase as well as amplitude and to observe the shape dynamics of the strip. The setup and results can be seen in Figure 7.

We then compare the experimental frequency response with predicted natural frequencies in Table 1. The natural frequency for a strip in the first mode can be estimated as:

$$f = \frac{1}{2L}\sqrt{\frac{T}{\mu}} \tag{1}$$

where L is the strip length, T the longitudinal tension and μ the linear density. We note that this approximation is adequate for small vibrations; however, for larger vibrations and elongations, more sophisticated models may be required [46]. There is a small discrepancy between measured and predicted natural frequencies that we attribute to the different clamps used in each setup and the permanent deformation that some samples may have undergone. In any case, it is possible to use the predicted values from Equation 1 as a reference for the resonant frequency of a strip for a given elongation (stretching).

Figure 7: Frequency response characterization. A) Setup for measuring frequency response. B) Strobe illumination fired when the sample is at the top displacement. C) Relative deformation of the different samples (in %) as a function of excitation frequency (in Hz).

Sample	Length at rest (± 1mm)	Length at measurement (± 1mm)	Resonance Frequency (Hz)	Damping Ratio (± 0.01)	Natural Frequency measured (Hz)	Natural Frequency predicted (Hz)	Difference Frequency measured vs. predicted
Elastane A	165	185	25.5 ± 0.5	0.118	25.9 ± 0.6	23.2 ± 0.3	2.7
Elastane B	168	185	28.0 ± 0.5	0.116	28.4 ± 0.5	26.7 ± 0.4	1.7
Elastane C	175	185	14 ± 0.5	0.107	14.2 ± 0.5	14.4 ± 0.2	0.2
Silicone	171	180	18.0 ± 0.5	0.083	18.1 ± 0.5	16.6 ± 0.3	1.5
Band	170	171	21.5 ± 0.5	0.105	21.7 ± 0.5	22.8 ± 0.7	1.1
Screen	160	183	27.5 ± 0.5	0.127	28.0 ± 0.6	23.9 ± 0.4	4.1

Table 1: Frequency response parameters per sample.

4.4 Plastic Deformations

Materials can have an elastic deformation, meaning that they will recover their shape and mechanical characteristics. However, beyond a certain elongation, they will suffer plastic deformations (or even breakage), meaning that they do not recover their original shape and mechanical properties.

The samples were subject to longitudinal stretching and relaxing 2 times while their force was measured, resetting the initial length after the first cycle. The permanent deformation and hysteresis can be seen in the offset between the initial and final point of each cycle in the elongation axis. The *Elastic band* (Figure 8 - Elastic Band) and *Silicone* were robust and their stiffness did not change significantly during the process. On the contrary, *Elastanes* and *Projection Screen* (Figure 8 - Projection Screen) suffered significant hysteresis; that would eventually lead to permanent deformation and even breakage (as seen in the discontinuities in Figure 6 - B).

The larger the plastic deformation, the more likely the end holes of the strips extend or break. However, in our setups, the holes were used to align the strips and set their length, and a clamping mechanism was used to avoid their deformation.

4.5 Material Selection

A good material for a FlexiVol diffuser should have the following characteristics:

Visual quality. It should provide an even diffusion of the light along the viewing angles but without blurring too much the patterns or attenuating them. The backprojected image should not distort

Figure 8: Stiffness of elastic band (blue dashed) and projection screen (purple dotted) as they are stretched and relaxed twice.

when the material is stretched, otherwise the volumetric graphics will look blurry at the top and bottom.

Stiffness. The force should not exceed 8 N when the material is stretched 80 mm – which is our maximum display volume height. Larger forces can damage the mechanical amplifiers and were found disturbing for the user. The selected samples fulfil this criteria, other materials were discarded given that they were too stiff and generated too much force on the user finger being unpleasant or even painful to touch when they were oscillating.

Frequency response. It may not be desirable to have the diffuser strip resonating at the operating frequency of the display because it would lead to larger forces on the display and large amplitudes provoking graphic distortions, harder to correct (see Section 5).

Plastic Deformation. It should have elastic deformation up to the elongation range. That is, the deformations should not be permanent since that would make the strips hang and give them different amplitudes when oscillating, thus incurring in discontinuous graphics in the rendering volume.

Given those criteria, **we selected the elastic bands** because they retain their mechanical and optical properties during elongations. Their visual quality is not the best since it has dark stripes along them, but when oscillating in a volumetric display that was not a significant problem. Elastane or the projection screen offered the best visual quality, yet they were permanently deformed even with small elongations and tend to have a natural frequency close to the operating frequency of the display, leading to large distortions. Silicone was elastic along the desired range but it did not diffuse light correctly. Adding ink or scattering particles to the mixture could improve its optical qualities; but in any case, its mechanical properties are not adequate since it is prone to rotate and not laying flat when oscillating – being an isotropic material in terms of stiffness.

The elastic bands were highly anisotropic; that is, they were elastic along their length but stiff on the other direction. This made them lay flat when oscillating on the display, passing around the finger but recovering their horizontal orientation when the finger was removed.

4.6 SVD Noise with Elastic Diffusers

We measured the noise levels using both the original rigid diffuser and the chosen elastic material (elastic band) on the Voxon V1 as a reference; and with and without the hand inside the display volume. Noise levels were captured 5 times using a Meterk MK09 with the sonometer's tip at 60 cm from the center of the display volume, which is the typical distance between the display and the user's ear. The noise level was 57.9 dB (std = 0.2) for the rigid diffuser (without safety dome) and 51.9 dB (std = 0.2) for the elastic diffuser. Inserting the hand inside the volume sligthly increased the noise level 53.0 dB (std = 0.3). Noise can be considered a relatively small issue and the noise levels actually decrease slightly by using an elastic diffuser.

5 Distortion Correction

The diffuser has to be elastic so that it deforms upon user's contact and does not damage their hand or the display. However, this same desired property makes the elastic diffuser become non-planar when it oscillates and thus the projected graphics look distorted (Figure 9).


Figure 9: Projection on a swept volumetric display. A) in a flat rigid diffuser, the cube is rendered correctly. B) the same projection on an elastic diffuser, renders the cube distorted.

Since the oscillation of the diffuser has a periodic motion, the shape of a planar diffuser and an elastic diffuser can be predicted at each timestamp. The inverse transformation from the elastic to the rigid diffuser can be applied vertex-wise to the rendered meshes to correct the distortion.

We assume that the deformation occurs only along the Z-axis (along the projection direction, i.e, top-bottom). Deformations along the Y-axis (back-forth) occur on a membrane diffuser which is continuous along the XY plane, but not significantly on individual strips (aligned along the X direction). Deformations along the X direction (left-right) occur on the strip physical material but not on the positions of the projected graphics.

We also assume that the shape of the elastic diffuser can be approximated as a sinusoid along the X direction as it oscillates. To simplify the fitting, we employed a parabole with 2 points at the sides of their holding frame and a central point at the center of the strip. Figure 10 shows the elastic diffuser at 3 different positions and their matching paraboles.


Figure 10: Top) Strips oscillating in the volumetric display. Left to right shows 3 different moments. Bottom) a parabole was fitted to the first strip.

We recorded the oscillating diffuser at 900 FPS with a Phantom MIRO-C211-8GB-M high speed camera for at least half a period of oscillation. For each captured frame, starting from the highest point of the frame to the lowest, a parabole was fitted to the strip. We processed only the descent of the diffuser to avoid the ambiguity of having 2 paraboles for each height. The correction was a fitted degree 6 polynomial

$$height = F(X, Z)$$

which given an X and Z position checks which parabole passes through that point and returns the height of the holding frame (Z of the frame). By applying this polynomial to the vertices of the meshes to be rendered, the distortion was significantly corrected. The effect of the distortion correction is shown in Figure 11. The display volume increases given the fabric deformation, but when using the correction, the effective volume remains the same as when using a rigid diffuser.

Figure 11: A calibration cube rendered on FlexiVol: left cube is uncorrected, right cube is corrected. A) viewed from the top-right corner. B) frontal view.

6 Reach-Through Interaction Design Space

FlexiVol is designed to enable *direct interaction*, allowing the users' hands to **reach-through** its elastic diffuser and interact with 3D content. In this section, we describe the interaction design space of FlexiVol, through Input and Interaction Techniques.

We draw inspiration from user-defined gestures with 2D elastic displays [57], touchscreens and AR [60], as well as virtual reality [27] to enhance the existing interaction design spaces for volumetric displays [17].

6.1 Input

Input represent the physical interaction from the user, and is here described in two categories: the *Interaction Space* - e.g., the *where* within FlexiVol; and the *Interaction Direction* - e.g., transversal or lateral interaction within FlexiVol's strips. We focus on users' interaction with hands (one or multiple fingers).

Figure 12: Frames extracted from high speed recordings of direct interactions with FlexiVol. A) 1. Interaction with the diffuser strips; 2. between the diffuser strips. B) Pinching gesture used to grab, scale or rotate objects. C) A sliding gesture initiated at a corner and directed towards the center.

6.1.1 Interaction Space. When the user hand reaches-through the display volume, it either collides with the diffuser strips, or passes between them.

Interaction with the diffuser strips. The diffuser deforms under the input type constraints, but the applied deformation is local to the colliding strip (see Figure 12 - A1). Strips can be independently or simultaneously interacted with; without altering the other ones.

Interaction between the diffuser strips. When interacting through the volume, the input type can also be moving between the diffuser strips (see Figure 12 - A2, B). In this case, the strips keep oscillating around the "obstacle".

However, when the user performs an interaction technique, most often at least one strip will be touched, providing a tangible dimension to the interaction. Even though the user's finger is between strips, if the finger moves, it will interact with a strip.

Interacting between strips can also be performed, for instance if the user goes below the diffuser and opens its fingers. This is however not recommended as it would cause the diffuser to pass through an area without visual rendering, hindering the interaction with 3D content.

Interaction with other diffuser types. We also tested interacting with a continuous diffuser (i.e., a membrane) instead of strips. This however causes the physical and visual deformations to be global over the whole surface; thus requiring more complex simulations of the surface deformation under the given number of contact points; as well as harder to calibrate distortion correction schemes. However, when used, the hand tracking does not suffer from the occlusion of the strips.

We considered the interaction space as a function of strips width and spacing. The wider the strips, the closer it is to a membrane diffuser and harder to visually correct for deformation. We also

Figure 13: Reach-through direct interaction techniques for FlexiVol: *Gestures* **and Actions**
.

tried strings and the visuals were good but the tactile feedback was similar to whipping the fingers, and was discarded. The wider the spacing between the strips, the less continuity there is in the visuals as the diffuser is not continuous. We also tested overlapping strips which resulted in worse visuals as the projection light had to project over two diffusers.

6.1.2 Interaction Direction. The input over the diffuser can be performed in any direction. The strips are indeed placed along a single longitudinal direction, but interactions can also occur along the lateral or diagonal directions (see Figure 12 - C). A tactile discontinuity can be perceived between strips, but the user input type (e.g., finger, stylus) can perform continuous gestures without getting stuck between/under a strip. The high operating frequency of the volumetric displays moves the strips fast enough so that the input (e.g., finger, stylus) does not perceive their oscillation and the strips remain under the input type's constraints. The strips oscillated with similar amplitude and phase, thus limiting potential entanglement with the use hand or between them.

6.2 Interaction Techniques

We selected the interaction techniques that are closer to real-life ones. Techniques that use metaphors such as rays, clutching with the thumb or micro-gestures were discarded since they are different from how we interact in the real world. In this section, we list the main user-gestures from FlexiVol's design space, along with user actions. We illustrate the interaction techniques as combinations of actions and gestures (Figure 1 - B).

6.2.1 Actions. We represent the main reach-through actions that users could perform in FlexiVol in Figure 13 - Actions), group by *Tracing, Selecting, Scaling, Rotating, Editing, Displacing, Navigating*.

Other examples of actions from elastic projection screens [57] can be applied to FlexiVol, such as pushing, twisting (e.g. as a combination of selecting, editing and rotating), squeezing (e.g., as a combination of selecting, editing and scaling down), or zooming-in/out (e.g., as scaling).

6.2.2 Gestures. We represent in Figure 13 the main user gestures in FlexiVol. They are a combination of gestures from the literature involving either touchscreens, multi-touch displays, and augmented/virtual reality 3D object manipulation [27, 55, 60]. In Figure 13, a common mapping from *Gestures* to *Actions* is shown. However, other mappings are possible. For instance, Selecting could be performed by *Touching* (as in our user study - see Section 7) or by *Grasping* between the fingers [12]. Displacing can be performed by

Pushing (as in our suggested applications - Section 9), or by *Grasping, and Pinching*. We also depict a walking Gesture for Navigating: *Alternate Swiping*. This gesture is metaphor-like and users can move their fingers to mimic a mini-character walking inside the display. When the ring and pinky finger are flexed the action is triggered, taking the direction of the backhand and the displacement of the movement of the index and heart fingers.

Different Input types can be used for each gesture. For instance, a pinching gesture can be performed with one hand (e.g., between two fingers of the same hand) or between two fingers of independent hands, to represent a *scaling-down* action. Similarly, swiping can involve one finger (e.g., as per the illustration), or two (e.g., as per gesture with Apple's trackpad to represent dragging actions [55]).

7 User Study

We conducted a user study to evaluate FlexiVol's usability and compare its direct interaction with a traditional volumetric display input modality (i.e., 3D mouse). We chose to reflect on this usability through both quantitative feedback: by comparing task completion time, accuracy and subjective cognitive workload; and qualitative feedback: regarding its use, first impressions, and envisioned applications involving FlexiVol. The comparison of touch direct interaction with mouse was performed for 2D screens [14], showing that the latter was more appropriate for single-point interaction. However, results may not be the same in true 3d environments.

We considered comparing FlexiVol with direct interaction on multi-touch screens or reach-through interaction in mid-air screens [38], however these devices do not render true 3d content coherently overlapped with the user hand.

The user study was split in **three** tasks (**Selection, Tracing** and **Docking**, see Figure 14), to evaluate the validity of part of the FlexiVol's design space through basic interactive tasks. This study was approved by the ethics, data protection and biosecurity committee of the university (PI-011/22) and conducted according to the pertinent legislation.

Figure 14: User study. A1) Selection task using the 3D mouse input modality. A2) Selection task using reach-through input modality. B) Docking task. C) Tracing task.

7.1 Participants

We recruited 18 participants (self-reported gender, 10 males and 8 females), age 20 to 40 years (mean = 26, std = 4.5). All participants were right-handed, two of them (11%) had previous experience with a 3D mouse, few of them were gaming-enthusiasts and were familiar with joystick interaction. Five of them (28%) had already visualized content on a volumetric display but never had interacted with it through a 3D mouse.

7.2 Apparatus

We evaluated two input modalities (Reach-through and 3D Mouse) using a custom diffuser and frame on a Voxon VX1 (as per the schematic Figure 2 - C). The diffuser was made of 7 strips placed with a stretching of 4 mm, the material was Elastic Band described in Section 4. The three tasks were designed on Unity3D with the Voxon SDK plugin. We applied a correction on the 3D meshes to correct the distortion due to the diffuser deformation upon oscillation (see Section 5).

For the 3D Mouse condition, we selected a *3D Connexion Space-mouse Compact*. For the Reach-through condition, we tracked the participant's index finger using two Papalook PA452 RGB cameras and Google's *Mediapipe* solution, using the library HandPose3D. The 3D coordinates of the user's index finger were computed and sent over UDP to the Unity application running on the display.

7.3 Overall Procedure

Participants entered the room and were informed about how to interact with the volumetric display, they saw the diffuser and were told how it would oscillate fast to make them see 3D objects. They were explained that the aim of the study was to evaluate conditions to interact with the volumetric graphics, with a 3D mouse and with their hands. The tasks were described to them and then they signed a consent form. During the three tasks, users were standing up in front of the volumetric display. The 3D mouse was placed on the same table as the volumetric display, within arm reach. Neither users' body nor head were constrained, though the 3D mouse was placed on the same table as the volumetric display. As per FlexiVol's design space, participants interacted both with the strips and between them, using lateral and transversal displacement across FlexiVol's volume.

After performing each Task (Selection, Tracing and Docking) with the two conditions (3D mouse and Reach-Through), users were asked to fill in an unweighted NASA-TLX questionnaire [20]. The duration of the study was in average 18 minutes (std = 2 min), and was followed by a semi-structured interview.

7.4 Design

We used a within-design and all users performed the three tasks with the two conditions (3D Mouse and Reach-Through). The conditions order was counterbalanced.

7.5 Task 1: Selection

We designed a first Selection task to evaluate how the input modality (reach-through vs. 3D mouse) impacts the completion time.

7.5.1 Description. Participants were first introduced to the interaction with the diffuser, and were asked to train by moving their index finger or the 3D mouse within the display volume. Once they felt confident, they had to touch a virtual button to start the experience. Red spheres of various sizes appeared in the display volume. They were asked to touch these spheres - either with the 3D Mouse or with their index finger. Once a sphere was touched, it disappeared and another sphere appeared. The sizes of the spheres were pseudo-randomized (within blocks) and the appearing position was randomized in the 3D space but constrained to not be overlapping with the disappearing one, to ensure participants had to move their cursor/hand for the next trial.

The 3D mouse and index finger cursor were represented by a 15 mm diameter sphere in the display volume.

7.5.2 Conditions. In this task, we controlled the Input modality (Reach-through; 3D Mouse) and the Sizes of the spheres. Three sizes were generated: small 15 mm diameter), medium (30 mm diameter) and large (60 mm diameter) in a pseudo-random order. Users were asked to perform 20 Blocks, giving 3 Sizes × 20 Blocks = 60 Trials per Input modality condition.

7.5.3 Measures. We measured the time for completion and recorded the sphere positions at each trial.

7.6 Task 2: Tracing

We designed a second task based on a tracing task, where the participants are required to follow a given 3D circuit.

7.6.1 Description. Participants were told to guide a ball 30 mm diameter) along a given circuit. They were then told how to interact in the display volume and could train in a practice scene. In the 3D Mouse condition, the 3D mouse was moving the ball was directly attached to the cursor. In the Reach-through condition, they were told they had to touch the ball; and when done, the ball would turn blue and follow the participant's finger. Once ready, participants told the evaluator that they could start. At each circuit, participants had to notify the researcher when they finished, the evaluator pressed the Space bar to start the next circuit.

7.6.2 Conditions. Five circuits were used with different characteristics (see Figure 16).

1. An elbow shape (one turn) with a displacement along XY.
2. A straight tube along XY.
3. A U-Shape (two turns) with a displacement along XY.
4. A descending tube with a displacement along XYZ (added depth compared to 2.).
5. A complex shape including 3 turns and a descending slope.

Participants traced the 5 Circuits with the two Input Modality conditions.

7.6.3 Measures. We recorded the cursor position (X, Y, Z) at all frames during the task. We also measured the time for completion and the cursor's travelled distance.

7.7 Task 3: Docking

We designed a third task to evaluate the effect of the Input modality on accuracy when manipulating 3D objects in the volumetric display, via a docking task.

7.7.1 Task description. A small white cube (*User cubes*, 30 mm side) appeared at pseudo-random position within the interaction space together with a larger red cube (*Target cubes*, 36 mm side). Participants were asked to align the white cube inside the red cube and to inform the researcher when they were done. They first were taught again how to interact with the two modalities and had a training scene to practice. Once again, in the 3D MOUSE condition, the 3D mouse was directly attached to the *User cube*. In the REACH-THROUGH condition, participants were asked to reach for the cube with their index finger, and when the contact occurred, the cube turned blue and was attached to the finger.

7.7.2 Conditions. We wanted to evaluate whether some directions were harder for docking than others. We thus evaluated 7 DIRECTIONS: X, Y, Z; XY, XZ, YZ; XYZ. The cubes were pseudo-randomly spawned within the DIRECTIONS condition, and this generation was constrained so that the cubes would not overlap. For instance, in the X direction, the Target cube would be generated at a random position; and the User cube would be generated with the same Y and Z coordinates but at a different random X position.

Participants performed 7 DIRECTIONS × 3 BLOCKS = 21 TRIALS per INPUT MODALITY.

7.7.3 Measures. We measured the time for completion at each trial, the 3D coordinates (X, Y, Z) of the target cube and the user cube at the end of the trial; and the cubes center-to-center distances.

8 Results

We conducted ANOVA test within the INPUT MODALITY condition and the other conditions (SIZE, DIRECTION) when the data allowed it (equal variances, normality). Posthocs pairwise T-tests with Bonferroni-corrected p-values were then computed.

8.1 Task 1: Selection

The Selection task was the fastest of the three tasks. There was both a significant effect on the INPUT MODALITY condition ($F_{(1,17)} = 44.5, p_{GG} < 0.001, \eta^2 = 0.47$) and the SIZE condition ($F_{(2,34)} = 39.24, p_{GG} < 0.001, \eta^2 = 0.07$).

8.1.1 Input Modality Effect. Results for the two modalities are displayed in Figure 15. Completion time was significantly smaller with Reach-through than with the 3D mouse ($p < 0.001$). In average, participants took 1.0 s (std = 0.5 s) to select spheres with Reach-through; whereas the 3D mouse selection was 1.9 s (std = 1.1 s).

8.1.2 Size Effect. As expected, completion time was significantly higher for small objects than medium ones ($p < 0.001$), medium and large ones ($p < 0.001$) and small and large ones ($p < 0.001$). Large objects were reached in average in 1.2 s (std = 0.7 s); medium ones in 1.4 s (std = 1.0 s); small ones in 1.6 s (std = 1.1 s).

8.1.3 Input Modality × Size Effect. We also noted a small but significant effect while crossing conditions ($F_{(2,34)} = 3.74, p_{GG} < 0.05, \eta^2 = 0.007$). Large targets were reached in average in 0.9 s (std

Figure 15: Boxplot of the Task Completion Time for *Task 1: Selection*. The reach-through input modality was significantly faster ($p < 0.001$) than the 3D Mouse.

= 0.3 s) in Reach-through, while it took 1.6 s (std = 0.8 s) with the 3D Mouse. In Reach-through, the medium targets were reached with similar performance as the large ones (mean = 1.0 s, std = 0.5 s), while the 3D mouse took 1.8 s (std = 1.1 s). Finally, the small objects were the hardest to reach, with an average of 1.2 s (std = 0.6 s) in Reach-through and 2.1 s (std = 1.3 s) with the 3D mouse.

8.2 Task 2: Tracing

We measured task completion time, and computed the overall travelled distance and position deviation in both modalities conditions.

8.2.1 Completion time. We did not note a significant different in time for completion using the two different Input modality techniques (in average 6.6 s, std = 2.9 s for Reach-through; 7.1 s, std = 2.6 s for 3D Mouse).

8.2.2 Travelled Distance. There was a small effect and significant difference on the cursor travelled distance within conditions ($F_{(1,15)} = 4.52, p_{GG} < 0.05, \eta^2 = 0.07$). The trajectories were shorter using the 3D mouse (average = 160 mm, std = 60 mm) than with Reach-through (average = 209 mm, std = 77 mm).

One cause could be that the 3D mouse was on a table, and the user could rest the wrist on the table as well, whereas the arm had no physical support during the reach-through interactions being more unstable and resulting in a more shaky cursor.

8.2.3 Position Deviation. We call position deviation the closest distance from the cursor to the virtual circuit (e.g. when the cursor is in contact with the circuit, the position deviation is null). We computed the average of all the trajectories along the circuit and averaged them out per circuit, user and input condition.

There was a small but significant effect on the position deviation ($F_{(1,17)} = 0.9, \eta^2 = 0.03$). In average, the position deviation with the Reach-through input modality is 3.8 mm (std = 3.2 mm); while the position deviation with the 3D mouse is in average 4.3 mm (std = 3.4 mm). The target circuits are displayed in Figure 16: their width is similar to the cursor ball's one. Results are relatively good for

a pure 3D task, we note that most of the positions are within the circuit's volume.
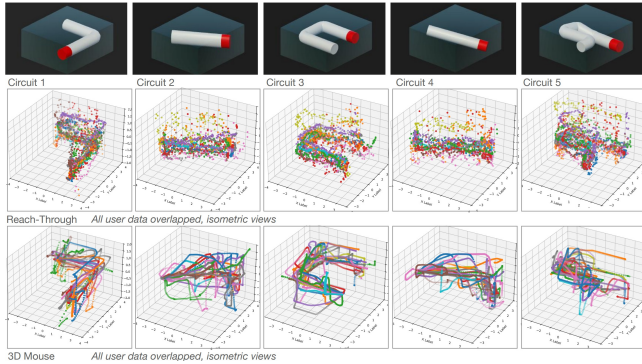
Figure 16: Target circuits and cursor position while tracing them in *Task 2: Tracing task*. All participants' trajectories (X, Y, Z) are shown for each INPUT MODALITY conditions.

8.3 Task 3: Docking

For the Docking task, we measured task completion time and accuracy. We computed accuracy as the distance center-to-center between the two cubes to align.

8.3.1 Completion Time. There was a significant difference in completion time between the two INPUT MODALITY ($F_{(1,16)} = 14.2, p_{GG} < 0.005, \eta^2 = 0.08$). The completion time with the Reach-through modality was in average 3.9 s (std = 1.3 s) while the 3D mouse one was 4.6 s per trial (std = 1.9 s). There was no effect on the DIRECTION for completion time or on the BLOCK condition.

8.3.2 Positional Error. Regarding positional error, there was a significant effect between REACH-THROUGH and 3D MOUSE ($F_{(1,15)} = 10.3, p_{GG} < 0.01, \eta^2 = 0.02$). The distances center-to-center were in average 22 mm (std = 23 mm) with the REACH-THROUGH, while the 3D MOUSE error in accuracy was 18 mm (std = 24 mm).

8.4 Mental Workload

Participants were asked to fill an unweighted NASA-TLX questionnaire [20] after each INPUT MODALITY condition.

We report the results in Figure 17. We noted a significant difference in mental demand, temporal demand and the overall score between the two conditions ($p < 0.05$). The average TLX-score for the 3D mouse was 31.2 (std = 13.5); for Reach-through was 22.8 (std = 12.1). Only physical demand was worse for "Reach-through" but not significantly. We discuss this result in the Discussion section (Section 8.6) to combine them with the Qualitative feedback (see Section 8.5) gathered during the semi-structured interviews.

8.5 Qualitative Feedback

We conducted semi-structured interviews after the participants finished the three tasks. The aim was to gather impressions on FlexiVol, mostly about usability and reluctance towards inserting their hand in the display, but also about enjoyment and potential future applications.
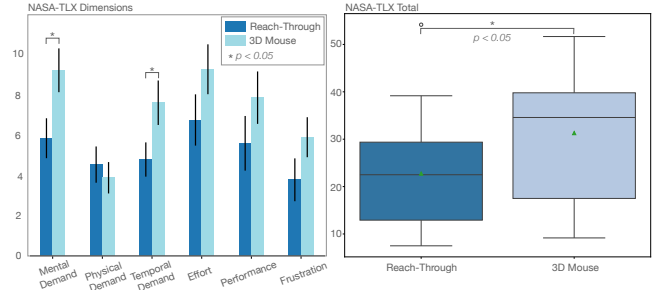
Figure 17: NASA-TLX scores. Left) Per dimension; Right) Overall score.

8.5.1 Preferred Condition and Confidence in Performances. Two participants (11%) preferred the 3D mouse condition, the rest preferred Reach-through (89%). More than half of the participants spontaneously commented that it felt "easier and more *natural* to interact with their fingers" (P2 mentioned it was "more spontaneous"). When asked about it, four participants mentioned "reaching-through was more intuitive" and two others added we are more "used-to interact with our hands". Most participants also mentioned that it was fun to reach-through inside the diffuser, and expressed this directly upon their first interaction with the diffuser.

Time for Completion. All participants but one (P16) mentioned they believed they had better performance on *time for completion* using their hands.

Accuracy. Six participants felt their performances regarding accuracy (e.g. in the Docking task) were better using the 3D mouse. Twelve mentioned they felt more confident when they performed the tasks with their hands (as per the NASA-TLX results, Figure 17 - Performance – the smaller the performance score, the more confidence in one's performance).

8.5.2 Reaching-Through apprehension. We asked participants what they felt when told they had to get their fingers inside the volume. The reactions were consistent as they all expressed their feelings followed by a "but". For instance, P8 thought reaching-through the diffuser "would hurt, *but it did not*". Similarly, P4 and P9 thought "the diffuser displacement would bother them, *but it did not*".

Almost all participants mentioned how *soft* the interaction felt (as opposed to the perception of hardness that they had at first). Three participants (P2, P3, P5) mentioned reaching-through FlexiVol was quite *pleasant*, and two participants (P3, P10) felt some tickles through the fabric vibrations.

8.5.3 Potential Use-cases and Interaction Techniques. Four participants noted they would have enjoyed trying more interaction techniques, such as pinching, rotating, scaling, or giving impulses with their fingers. Three other participants mentioned they wanted to put their whole hand inside the volume, and would have liked to catch and power grab a ball.

Regarding use-cases or applications, P5 mentioned he enjoyed not having to wear any VR headset and added that he would like to use FlexiVol for 3D modeling, 3D sketches and painting. Three participants mentioned they could picture using FlexiVol to visualize

3D plans and easily collaborate with their friends. Six participants imagined using FlexiVol for medical purposes, such as surgery - to study prior to it or for teleoperation. Videogames and chess games were also mentioned; P14 and P18 suggested film editing; to select, cut, and visualize 3D content, whether 2D or 3D; and P17 mentioned an online shopping simulator.

8.6 Discussion

We note that completion time is significantly smaller using reach-through than using the 3D mouse for all studies. We however note that a trade-off between completion time and accuracy.

8.6.1 Completion Time vs. Accuracy. We note a trade-off between *Completion time* and *Accuracy*, especially notable in *Task 3: Docking*. All participants but one did perceive that their completion time performances were better using the reach-through condition, and a majority also believed their accuracy was better using their hand; while the others mentioned the accuracy was higher using the mouse. Three main reasons can explain this: a. while participants enjoyed the proprioception of being *within* the diffuser, seeing it from the outside with an indirect interaction allows for a better overall view and a more accurate interaction; b. the recurring fat finger issue from touchscreens [23] also appears in reach-through 3D interaction and alters the accuracy performances; c. the tracking of the user's hand can also be improved using more cameras, therefore enhancing the overall accuracy of the task.

8.6.2 Mental Workload. Participants mentioned how easier it felt to use the hands, and used words such as "natural, intuitive, spontaneous, used-to it". These qualitative feedback are consistent with the NASA-TLX results - where participants felt more frustration, mental demand and effort using the 3D mouse, for less performance.

Regarding physical demand, it is coherent that using the reach-through interaction is more tiring (even though non-significant, see Figure 17 - Physical demand): the 3D mouse was on a table and user could rest the elbow or wrist on the table, whereas the finger interaction required to lift the arm to reach inside the display. Finding body postures or layouts (display – user) that are more ergonomic could mitigate this issue.

8.6.3 Body position influence on performance. We do not believe that the user's body and head position significantly influenced the performance in the 3D interaction tasks; apart from the physical demand. Participants could be more fatigued with the reach-through modality as directly interacting with the 3D content was more engaging. Regarding head position, participants looked at the display volume both from the top and sides to localize the 3D objects or during the tracing task; however they did so in both input modalities.

8.6.4 Future interaction techniques and use-cases. While participants could only point to 3D objects to either select or displace them in the user studies, pinching the object to grasp it (as per Figure 1 - B1) , rotate it (as per Figure 1 - B3) or scale it was mentioned multiple times. The use of more fingers, the full hand, or collaborating with other users was also mentioned. We were comforted about the potential of volumetric displays for 3D visualisation and manipulation, especially when P5 mentioned picturing himself using

FlexiVol instead of a VR headset which he considered heavy and intrusive when performing creative tasks such as 3D painting.

9 Potential Applications

Virtual pets (Figure 18.A) displayed in FlexiVol could generate extra attachment with the user. Directly caressing the creature to comfort it can provide stronger connections with a virtual entity. Other affective interactions are possible, like the pet holding your finger as you walk it around. Similarly, when purchasing on an online shop, objects that we carry or wear can be more desirable or we can picture better how we would take them with us.

We have shown basic interactions such as selection, manipulation and tracing in the user studies. These tasks can be combined to create more complex applications such as 3D scene editors (Figure 18.B). A single finger is used to select trees, pinching gestures to raise or lower the terrain, and a gesture of walking with the index and heart fingers to navigate through the map.

Direct edition of a mesh is shown in Figure 18.C, pinching gestures are used for direct vertex manipulation whereas pushing with the finger can modify surfaces. Ten years ago, a conceptual video [48] speculated with the potential of Claytronics showing direct interactions to model a mock-up car. Reach-through volumetric displays can be a platform to realize some of the concepts that were expected from Claytronics, Programmable matter, Radical Atoms [25] or similar paradigms.

Figure 18: Potential applications for FlexiVol. A) Petting a cute bee. B) Editing a 3D landscape. C) 3D model of a car being modified by an user translating one vertex. D) Data visualization and manipulation. E) File sorting.

Menus and Widgets can be adapted to FlexiVol. For example, the 2D menu on the iWatch (Apple Inc.) has a 2D layout of icons that shift and zooms in as the user moves the finger, this can be adapted to 3D. Also, tree directories can be navigated by reaching with the finger to the desired folder while the previous and next level are render above and below the finger. We exemplify a Data Viz application in Figure 18.D and a file transfer between folders in Figure 18.E.

10 Discussion

10.1 Limitations

Long-term fatigue test. We checked the robustness of the diffusers by stretching them several times, longer term fatigue and wear

tests should be conducted to characterise how many hours of user interaction the diffusers support without getting permanently deformed. In any case, we note that the elastic bands used in the user studies had no notable deformation and kept the same amplitude upon oscillation.

Other Diffuser types. We used out-of-the-shelf materials for the diffusers. We tried a composite of silicone with projection screen but the resulting material had issues such as detachments and too much stiffness. For a commercial FlexiVol device, advanced materials should be tested; e.g., adding ink or scattering particles to make silicone more optically diffuse and micro-patterning it for adding stiffness anisotropy.

Hand-Tracking capabilities. For hand-tracking, we used two stereo cameras with Google's Mediapipe library. It gave usable results, with better performance than Leap Motion, Meta Quest 2 or markers on the fingers which tend to get entangled with the strips and limit the come and interact paradigm of volumetric displays. Yet some interactions were limited by our current hand-tracking.

Ergonomics. Physical demand was reported to be slightly higher using direct interaction than with a traditional 3D mouse. FlexiVol layout can be made more ergonomic to avoid that the users have to hang their arms in mid-air. For example, integrating FlexiVol inside a desk so that users can reach inside it while resting their elbows.

Display Volume and Resolution. The display size (190x190x80 mm^3) and resolution (980x980x90 vx) can be improved to provide a better user experience. However, we think that the technical specifications of the current realization of FlexiVol permitted to conduct a user study and showcased the advantages of direct interaction on a volumetric display. Some future work discusses how to increase the display volume.

10.2 Future Work

Larger Sweeping Methods. Current commercial swept-volumetric displays are moving from oscillating planar diffusers to rotating helices [44]. A rotating diffuser can be larger and the mechanical part is heavily simplified. We reckon that it is possible to design elastic helical diffusers, for example in the shape of a central rotating axis with stripes attached to it. The strips are only attached on one side (to the axis) and due to centrifugal forces they get perpendicular to the axis, becoming a cylindrical projection volume.

Adaptative Rendering. We corrected the graphic distortion from the deformation of the elastic diffusers; however, the user hand also deforms the contacted strips. An adaptive rendering algorithm that adapts the projection in response to the user touch should be developed as future work.

Encountered-Haptics. When the users introduce their hand in FlexiVol there is tactile feedback (users described it as soft). However, the tactile feedback is provided throughout the whole volume. Haptics technologies can be added to FlexiVol so that only when the user hand encounters a virtual object, tactile stimuli are provided. Focused ultrasound can pass through some fabrics (depending on the fabric [7]) to create tactile stimuli on the fingertips, the ultrasonic array could be integrated with the display placed below the diffuser with a hole to let the light projection pass. By sewing conductive threads along the strips it would be possible to apply electrotactile feedback. Another option to explore is adding an overtone of 200 Hz on the strip oscillation to feel their vibration.

Augmented Reality. Only the user hand has been inserted in the display volume but other objects can be inserted to render graphics around them. In the same way that pens are used with touchscreens to provide more accuracy and mitigate the fat finger problem, these stylus can also be used for FlexiVol. Use cases of augmented reality can be developed, for example by rendering a watch or rings around the user hand, or by adding virtual gears to a physical piece that is inserted in the display volume.

Collaboration. Given the limitations in display size, collaborative applications were not studied. Co-located interaction in volumetric displays is an exciting direction to explore since its come-and-interact paradigm is especially suited for various users.

11 Conclusion

We have shown the concept of modifying swept volumetric display with an elastic diffuser to allow users to reach inside the display volume for interaction. When the elastic diffuser collides with the hand it does not damage the user and the localised deformation does not significantly distort the graphics. We showed in a user study how FlexiVol with direct interaction provides relatively good results in selection, docking and tracing tasks when compared with a 3D mouse using indirect interaction. FlexiVol provides coherent focus accommodation for both the real hand and the virtual graphics facilitating further the direct interaction between the user and the rendered objects. We believe that this simple yet significant improvement on volumetric displays creates new opportunities for exploring the unique advantages of volumetric displays and direct reach-through interaction.

Acknowledgments

References

[1] 3dconnexion. 2021. Space Mouse 3D input Device. http://www.3dconnexion.fr/nc/company/press-room/
[2] Isayas Berhe Adhanom, Paul MacNeilage, and Eelke Folmer. 2023. Eye tracking in virtual reality: a broad review of applications and challenges. *Virtual Reality* 27, 2 (2023), 1481–1505.
[3] Ray Asahina, Takashi Nomoto, Takatoshi Yoshida, and Yoshihiro Watanabe. 2021. Realistic 3D swept-volume display with hidden-surface removal using physical materials. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 113–121.
[4] Ravin Balakrishnan, George W Fitzmaurice, and Gordon Kurtenbach. 2001. User interfaces for volumetric displays. *Computer* 34, 3 (2001), 37–45.
[5] Peter C. Barnum, Srinivasa G. Narasimhan, and Takeo Kanade. 2010. A multi-layered display with water drops. In *ACM SIGGRAPH 2010 papers*. ACM, Los Angeles California, 1–7. https://doi.org/10.1145/1833349.1778813
[6] Steve Bryson. 2005. Direct Manipulation in Virtual Reality. In *Visualization Handbook*. Elsevier, 413–430. https://doi.org/10.1016/B978-012387582-2/50023-X
[7] Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. 2013. UltraHaptics: multi-point mid-air haptic feedback for touch surfaces. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. 505–514.
[8] Alvaro Cassinelli and Masatoshi Ishikawa. 2005. Khronos projector. In *ACM SIGGRAPH 2005 Emerging technologies on - SIGGRAPH '05*. ACM Press, Los Angeles, California, 10. https://doi.org/10.1145/1187297.1187308
[9] Dhairya Dand and Robert Hemsley. 2013. Obake: interactions on a 2.5 D elastic display. In *Adjunct Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*. 109–110.

[10] Nguyen-Thong Dang. 2007. A survey and classification of 3D pointing techniques. In *2007 IEEE international conference on research, innovation and vision for the future*. IEEE, 71–80.

[11] S. Eitoku, T. Tanikawa, and Y. Suzuki. 2006. Display Composed of Water Drops for Filling Space with Materialized Virtual Three-dimensional Objects. In *IEEE Virtual Reality Conference (VR 2006)*. 159–166. https://doi.org/10.1109/VR.2006.51 ISSN: 2375-5334.

[12] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M. Dollar, and Danica Kragic. 2016. The GRASP Taxonomy of Human Grasp Types. *IEEE Transactions on Human-Machine Systems* 46, 1 (Feb. 2016), 66–77. https://doi.org/10.1109/THMS.2015.2470657

[13] Unai Javier Fernández, Iosune Sarasate, Iñigo Ezcurdia, Manuel Lopez-Amo, Ivan Fernández, and Asier Marzo. 2024. PointerVol: A Laser Pointer for Swept Volumetric Displays. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3654777.3676432

[14] Clifton Forlines, Daniel Wigdor, Chia Shen, and Ravin Balakrishnan. 2007. Direct-touch vs. mouse input for tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 647–656. https://doi.org/10.1145/1240624.1240726

[15] Tatsuki Fushimi, Asier Marzo, Bruce W Drinkwater, and Thomas L Hill. 2019. Acoustophoretic volumetric displays using a fast-moving levitated particle. *Applied Physics Letters* 115, 6 (2019).

[16] E Bruce Goldstein and James R Brockmole. 2002. *Sensation and perception*. Vol. 90. Wadsworth-Thomson Learning Pacific Grove, CA, USA.

[17] Tovi Grossman and Ravin Balakrishnan. 2006. The design and evaluation of selection techniques for 3D volumetric displays. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*. 3–12.

[18] Tovi Grossman, Daniel Wigdor, and Ravin Balakrishnan. 2004. Multi-finger gestural interaction with 3d volumetric displays. In *Proceedings of the 17th annual ACM symposium on User interface software and technology*. 61–70.

[19] Joonku Hahn, Woonchan Moon, Hosung Jeon, Minwoo Jung, Seongju Lee, Gunhee Lee, and Muhan Choi. 2023. Volumetric 3D Display: Features and Classification. *Current Optics and Photonics* 7, 6 (2023), 597–607.

[20] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Advances in Psychology*. Vol. 52. Elsevier, 139–183. https://doi.org/10.1016/S0166-4115(08)62386-9

[21] Otmar Hilliges, David Kim, Shahram Izadi, Malte Weiss, and Andrew Wilson. 2012. HoloDesk: direct 3d interactions with a situated see-through display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2421–2430.

[22] Ryuji Hirayama, Diego Martinez Plasencia, Nobuyuki Masuda, and Sriram Subramanian. 2019. A volumetric display for visual, tactile and audio presentation using acoustic trapping. *Nature* 575, 7782 (2019), 320–323.

[23] Paul Huber. 2015. Inaccurate input on touch devices relating to the fingertip. *Media Informatics Proseminar on" Interactive Surfaces* 31 (2015).

[24] Lumi Industries. 2017. VVD: Volumetric Visualization Device. https://www.lumindustries.com/3d-vis

[25] Hiroshi Ishii, Dávid Lakatos, Leonardo Bonanni, and Jean-Baptiste Labrune. 2012. Radical atoms: beyond tangible bits, toward transformable materials. *interactions* 19, 1 (2012), 38–51.

[26] Jacek Jankowski and Martin Hachet. 2013. A survey of interaction techniques for interactive 3D environments. In *Eurographics 2013-STAR*.

[27] Ying Jiang, Congyi Zhang, Hongbo Fu, Alberto Cannavò, Fabrizio Lamberti, Henry Y K Lau, and Wenping Wang. 2021. HandPainter - 3D Sketching in VR with Hand-based Physical Proxy. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 412, 13 pages. https://doi.org/10.1145/3411764.3445302

[28] Andrew Jones, Ian McDowall, Hideshi Yamada, Mark Bolas, and Paul Debevec. 2007. Rendering for an interactive 360° light field display. *ACM Transactions on Graphics* 26, 3 (July 2007), 40. https://doi.org/10.1145/1276377.1276427

[29] Abhijit Karnik, Archie Henderson, Andrew Dean, Howard Pang, Thomas Campbell, Satoshi Sakurai, Guido Herrmann, Shahram Izadi, Yoshifumi Kitamura, and Sriram Subramanian. 2011. Vortex: Design and implementation of an interactive volumetric display. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. 2017–2022.

[30] Hidei Kimura, Akira Asano, Issei Fujishiro, Ayaka Nakatani, and Hayato Watanabe. 2011. True 3D display. In *ACM SIGGRAPH 2011 Emerging Technologies*. 1–1.

[31] Philip Kingsley, J Rossiter, and S Subramanian. 2012. eTable: A haptic elastic table for 3D multi-touch interactions. *Master's thesis. University of Bristol* (2012).

[32] Kota Kumagai, Shun Miura, and Yoshio Hayasaki. 2021. Colour volumetric display based on holographic-laser-excited graphics using drawing space separation. *Scientific Reports* 11, 1 (Nov. 2021), 22728. https://doi.org/10.1038/s41598-021-02107-3

[33] Miu-Ling Lam, Bin Chen, and Yaozhung Huang. 2015. A novel volumetric display using fog emitter matrix. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. 4452–4457. https://doi.org/10.1109/ICRA.2015.7139815 ISSN: 1050-4729.

[34] LegenDay. 2024. Latex vs Silicone: A Comprehensive Comparison. https://legenday.com.cn/latex-vs-silicone/

[35] Ivan T Lima and Val R Marinov. 2010. Volumetric display based on two-photon absorption in quantum dot dispersions. *Journal of Display Technology* 6, 6 (2010), 221–228.

[36] Andrew Maimone, Andreas Georgiou, and Joel S Kollin. 2017. Holographic near-eye displays for virtual and augmented reality. *ACM Transactions on Graphics (Tog)* 36, 4 (2017), 1–16.

[37] Dan Maloney. 2023. A Volumetric Display With A Star Wars Look And Feel. https://hackaday.com/2023/06/20/a-volumetric-display-with-a-star-wars-look-and-feel/

[38] Yasuaki Monnai, Keisuke Hasegawa, Masahiro Fujiwara, Kazuma Yoshino, Seki Inoue, and Hiroyuki Shinoda. 2014. HaptoMime: mid-air haptic interaction with a floating virtual screen. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) *(UIST '14)*. Association for Computing Machinery, New York, NY, USA, 663–667. https://doi.org/10.1145/2642918.2647407

[39] Mathias Müller, Anja Knöfel, Thomas Gründer, Ingmar Franke, and Rainer Groh. 2014. FlexiWall: Exploring Layered Data with Elastic Displays. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*. Association for Computing Machinery, New York, NY, USA, 439–442. https://doi.org/10.1145/2669485.2669529

[40] Yoichi Ochiai, Kota Kumagai, Takayuki Hoshi, Jun Rekimoto, Satoshi Hasegawa, and Yoshio Hayasaki. 2016. Fairy lights in femtoseconds: Aerial and volumetric graphics rendered by focused femtosecond laser combined with computational holographic fields. *ACM Transactions on Graphics (TOG)* 35, 2 (2016), 1–14.

[41] Jae-Hyeung Park and Byoungho Lee. 2022. Holographic techniques for augmented reality and virtual reality near-eye displays. *Light: Advanced Manufacturing* 3, 1 (2022), 137–150.

[42] Shreya K Patel, Jian Cao, and Alexander R Lippert. 2017. A volumetric three-dimensional digital light photoactivatable dye display. *Nature communications* 8, 1 (2017), 15239.

[43] Ken Perlin. [n. d.]. HoloDust. https://mrl.cs.nyu.edu/~perlin/experiments/holodust/

[44] Voxon Photonics. 2021. The Voxon VX1 Volumetric Display. https://voxon.co/voxon-vx1-available-for-purchase/

[45] Ismo Rakkolainen, Tobias Höllerer, Stephen DiVerdi, and Alex Olwal. 2009. Mid-air display experiments to create novel user interfaces. *Multimedia tools and applications* 44 (2009), 389–405.

[46] Singiresu S Rao. 2019. *Vibration of continuous systems*. John Wiley & Sons.

[47] Stephan Reichelt, Ralf Häussler, Gerald Fütterer, and Norbert Leister. 2010. Depth cues in human visual perception and their realization in 3D displays. In *Three-Dimensional Imaging, Visualization, and Display 2010 and Display Technologies and Applications for Defense, Security, and Avionics IV*, Vol. 7690. SpIE, 92–103.

[48] David S. Ricketts. 2014. Claytronics: Building Matter from Microscale Robots. https://www.youtube.com/watch?v=-qw0hjeQ0Yw

[49] Deepak Ranjan Sahoo, Kasper Hornbæk, and Sriram Subramanian. 2016. Tablehop: An actuated fabric display using transparent electrodes. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 3767–3780.

[50] Hideo Saito, Hidei Kimura, Satoru Shimada, Takeshi Naemura, and Jun Kayahara. 2008. Laser-plasma scanning 3D display for putting digital contents in free space National Institute of Advanced Industrial Science and Technology. *Spie-Is&T* 6803, 68030-1 (2008), 1–10.

[51] Lina Sawalha, Monte P Tull, Matthew B Gately, James J Sluss, Mark Yeary, and Ronald D Barnes. 2012. A large 3D swept-volume video display. *Journal of Display Technology* 8, 5 (2012), 256–268.

[52] Takashi Shibata, Joohwan Kim, David M Hoffman, and Martin S Banks. 2011. The zone of comfort: Predicting visual discomfort with stereo displays. *Journal of vision* 11, 8 (2011), 11–11.

[53] DE Smalley, E Nygaard, K Squire, J Van Wagoner, J Rasmussen, S Gneiting, K Qaderi, J Goodsell, W Rogers, M Lindsey, et al. 2018. A photophoretic-trap volumetric display. *Nature* 553, 7689 (2018), 486–490.

[54] Daniel Smalley, Ting-Chung Poon, Hongyue Gao, Joshua Kvavle, and Kamran Qaderi. 2018. Volumetric displays: turning 3-D inside-out. *Opt. Photonics News* 29, 6 (2018), 26–33.

[55] Hemant Bhaskar Surale, Aakar Gupta, Mark Hancock, and Daniel Vogel. 2019. Tabletinvr: Exploring the design space for using a multi-touch tablet in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.

[56] Yutaka Tokuda, Mohd Adili Norasikin, Sriram Subramanian, and Diego Martinez Plasencia. 2017. MistForm: Adaptive shape changing fog screens. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 4383–4395.

[57] Giovanni Maria Troiano, Esben Warming Pedersen, and Kasper Hornbæk. 2014. User-defined gestures for elastic, deformable displays. In *Proceedings of the 2014 international working conference on advanced visual interfaces*. 1–8.

[58] Yoshihiro Watanabe, Alvaro Cassinelli, Takashi Komuro, and Masatoshi Ishikawa. 2008. The deformable workspace: A membrane between real and virtual space. In *2008 3rd IEEE International Workshop on Horizontal Interactive Human Computer Systems*. IEEE, 145–152.

[59] Kyungwon Yun, JunBong Song, Keehong Youn, Sungmin Cho, and Hyunwoo Bang. 2013. ElaScreen: exploring multi-dimensional data using elastic screen. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 1311–1316.

[60] Lixiang Zhao, Tobias Isenberg, Fuqi Xie, Hai-Ning Liang, and Lingyun Yu. 2025. SpatialTouch: Exploring Spatial Data Visualizations in Cross-reality. *IEEE Transactions on Visualization and Computer Graphics* (2025).

A Appendix

A.1 Custom-Made SVD

Figure 19: Custom-made volumetric display (hardware parts are described in Section 3 and available on this repository.
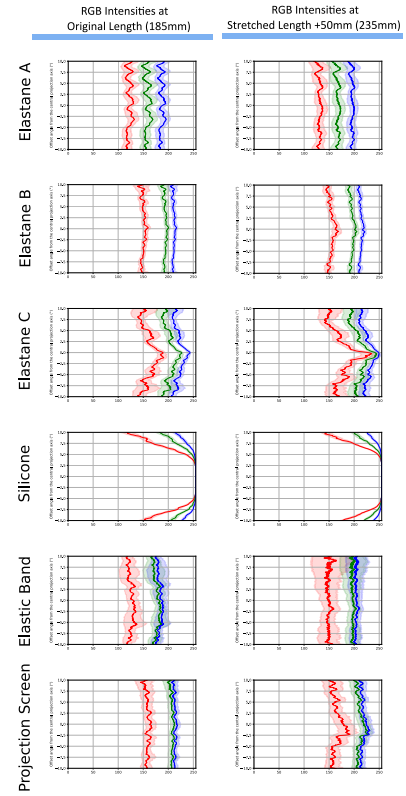
A.2 Optical Diffusion

Figure 20: Vertical cross-sectional analysis (1 px width) illustrating the average and 95-CI Savitzky-Golay filtered intensity levels of red, green and blue intensities at different offset angles from the central projection axis for the strips.